Eye-Movement behavior identification for AD diagnosis

Juan Biondi *[1,2], Gerardo Fernandez[1], Silvia Castro[2], and Osvaldo Agamennoni[1,3]

[1]Instituto de Investigaciones en Ingeniería Eléctrica (IIIE), Departamento de Ingeniería Eléctrica y de Computadoras (DIEC), Universidad Nacional del Sur (UNS) - CONICET, Bahía Blanca, Argentina
[2]Laboratorio de Visualización y Computación Gráfica (VyGLab), Departamento de Ciencias e Ingeniería de la Computación (DCIC), Universidad Nacional del Sur (UNS), Bahía Blanca, Argentina
[3]Comisión de Investigaciones Científicas de la Provincia de Buenos Aires (CIC), Argentina

arXiv:1702.00837v2 [cs.NE] 18 Feb 2017

Abstract

In the present work, we develop a deep-learning approach for differentiating the eye-movement behavior of people with neurodegenerative diseases over healthy control subjects during reading well-defined sentences. We define an information compaction of the eye-tracking data of subjects without and with probable Alzheimer's disease when reading a set of well-defined, previously validated, sentences including high-, low-predictable sentences, and proverbs. Using this information we train a set of denoising sparse-autoencoders and build a deep neural network with these and a softmax classifier. Our results are very promising and show that these models may help to understand the dynamics of eye movement behavior and its relationship with underlying neuropsychological correlates.

Keywords: Eye-tracking, Deep-learning, Alzheimer Disease

1 Introduction

Alzheimers disease (AD) is a nonreversible neurodegenerative disease characterized by progressive impairment of cognitive and memory functions that develops over a period of years being the most prevalent cause of dementia in elderly subjects. Initially, people experience memory loss and confusion, which may be mistaken for the kinds of memory changes that are sometimes associated with normal aging Waldemar et al. (2007). The subtle changes in behavior and response of the early manifestation of this disease difficult the diagnose by using the classical neuropsychological tests such as minimental. The use of more advanced diagnosis tools such as MRI and PET results critical for its early diagnose. Since AD is nonreversible, its early treatment can improve the patient's life delaying the full manifestation of the disease. In the last years, the study of the eye movement, known as *eye-tracking*, during reading, has proved to help performing this task (Fernández et al., 2015b, 2016b, 2015a, 2013).

Reading is a cognitive activity that has received considerable attention of researchers to evaluate human cognitive performance. This requires the integration of several central cognitive subsystems, from attention and oculomotor control to word identification and language comprehension. Eye movements show a reproducible pattern during normal reading. Each eye movement ends up in a fixation point, which allows the brain to process the incoming information and to program the following saccade. Different neuropsychiatric pathologies produce abnormalities in eye movements and disturbances in reading, having each of them a particular pattern that can be registered and measured (Fernández et al., 2016c,a; Holzman et al., 1974; Iacono et al., 1992; Riby and Hancock, 2009; Kellough et al., 2008). Eye movements can be classified into three groups:

- Movement for maintaining the image on the fovea (area of the retina with acuity vision), compensating head or object movements;

*Corresponding Author.
Email: *juan.biondi@uns.edu.ar*

- Movements for shifting the eyes, when the attention changes from one object to another. There are subtypes of shifting movements: saccades (looking for a new center of visual attention), monitoring and vergence (slower than saccades and are responsible for carrying the image of interest to both foveae, allowing stereoscopic vision);
- Movements of binocular fixation that also prevent fading of the image. These movements have three variations: tremor, drift, and microsaccade.

Saccades are rapid big eye movements particularly important from the cognitive point of view since cognitive processes have a direct influence on such movements. Each saccade has its direction. People, depending on language, read from left to right and most of the saccadic eye movements are oriented accordingly. These normal reading movements are called forward saccades. Reading movements going from right to left are called regressions. The saccade movement alternates with a fixation made when the eyes are directed to a particular target (See (Rayner, 1998) for a review). As shown in (Fernández et al., 2013), patients with early Alzheimer disease show alterations during the execution of tasks, such as reading, and these alterations can be related to an impairment in their working memory (Fernández et al., 2014a, 2016b). In fact, it has been shown that through this differentiation in the eye-movements, it is possible to infer a diagnosis (Fernández et al., 2015a).

The use of computer-aided diagnosis is a key challenge since the growth of computational power permits the creation of more complex models. These models can be used to create biomarkers that help in disease identification. Since the popularization of *deep-learning* neural networks Schmidhuber (2015); Deng and Yu (2014), many efforts have been made in their use in the field of medicine. This technique is commonly used in conjunction with imaging diagnoses such as PET or MRI mainly because the feature representation that this technology provides may help even when data is incomplete Li et al. (2014). Specifically, there were advances in the detection and pattern differentiation of the physical brain alterations that neurodegenerative diseases produce, such as AD and mild cognitive impairment (MCI) Suk and Shen (2013); Suk et al. (2014). Even there were advances in its early diagnostic Liu et al. (2014). The problem is that, when a brain physical alteration is observable, the damages made to the brain are irreversible (even though the disease is in an early stage) and may cause deterioration in the quality of life of the patient. The eye-tracking technique allows us to find subtler changes that were made by the brain to alleviate small memory deficits in the patient. These changes are not noticed by the patient but small changes in the way they read our set of sentences can be found with the technique presented in this paper.

In this work, we use a deep-learning neural network trained on reading information extracted from controls and patients with probable AD in order to identify the patterns made by them in the reading process and later cluster them in their respective groups. Throughout this work, we may use AD patients and patients with probable AD in an interchangeable manner because of the nature of the AD diagnosis. The hypothesis was that using deep-learning in the feature identification of the key characteristics of the patient's eye behavior during reading sentences may lead to a correct classification that can be used to infer a diagnosis. Using this type of technology may improve the results obtained in (Fernández et al., 2015a) since it provides a smaller granularity in the detection of the disease and consequently, a better performance. Additionally, this technology allows us to improve the effectiveness of the classification as we collect more ground truth subjects.

2 Methods

2.1 Ethics Statement

The investigation adhered to the principles of the Declaration of Helsinki and was approved by the Institutional Bioethics Committee of the Hospital Municipal de Agudos (Bahía Blanca, Buenos Aires, Argentina). All patients and their caregivers, and all control subjects signed an informed consent prior to their inclusion in the study.

2.2 Participants

Twenty patients (12 females and 8 males; mean age 69 years, $SD = 7.3$ years) with the diagnosis of probable AD were recruited at Hospital Municipal of Bahía Blanca (Buenos Aires, Argentina). The

clinical criteria to diagnose AD at its early stages remains under debate McKhann et al. (1984). In the present work, the diagnosis was based on the criteria for dementia outlined in the Diagnostic and Statistical Manual of Mental Disorders (DSM-IV). All AD patients underwent a detailed clinical history revision, physical/neurological examination and thyroid function test. They all presented an APO E3E4 Genotype. Magnetic resonance images were obtained from twelve patients and computerized tomography scans from the other eight. All the patients underwent biochemical analysis to discard other common pathologies (hemoglobin, full blood count, erythrocyte sedimentation rate, urea and electrolytes, blood glucose). All this data provided a more precise diagnosis of AD. Patients were excluded if: (1) they suffered from any medical conditions that could account for, or interfere with, their cognitive decline; (2) had evidence of vascular lesions in computed tomography or FMRI; (3) had evidence for an Axis I diagnosis (e.g. major depression or drug abuse) as defined by the DSM-IV. To be eligible for the study, patients had to have at least one caregiver providing regular care and support. Patients taking cholinesterase inhibitors (ChE-I) were not included. None of the subjects were taking hypnotics, sedative drugs or major tranquilizers. The control group consisted of 40 elderly adults (29 female and 11 male); mean age 71 years ($SD = 6.1$), with no known neurological or psychiatric disease according to their medical records, and no evidence of cognitive decline or impairment in daily activities. A one-way ANOVA showed no significant differences between the ages of AD and Control individuals. Those participants diagnosed of suffering from Ophthalmologic disease such as glaucoma, visually significant cataract or macular degeneration as well as visual acuity less than 20/20 were excluded from the study.

The mean scores of Controls and AD patients in the Mini-Mental State Examination (MMSE) Folstein et al. (1975) were 27.8 ($SD = 1.0$) and 24.2 ($SD = 0.8$), respectively, the latter suggesting early mental impairment. A one-way ANOVA evidenced significant differences between MMSE in AD patients and Controls ($p < 0.001$). The mean score of AD patients in the Adenbrooks Cognitive Examination - Revised (ACE-R) Mioshi et al. (2006) was 84.4, ($SD = 1.1$), the cut-off being of 86. The mean school education trajectories in AD patients and Controls were 15.2 ($SD = 1.3$) years and 15.1 ($SD = 1.0$) years, respectively. A one-way ANOVA showed no significant differences between the education of AD and Control individuals.

2.3 Apparatus and eye movement data

Single sentences were presented on the center line of a 20-inch LCD Monitor ($1024x768$ pixels resolution; font: regular; New Courier; 12 points, 0.2° in height). Participants sat at a distance of 60 cm from the monitor. Head movements were minimized using a chin rest. Correction for the 60 cm viewing distance was performed by using the *Eyelink1000* corneal reflection system, which assessed changes in gaze position by measuring both the reflection of an infrared illuminator on the cornea and the pupil size, by means of a video camera sensitive to light in the infrared spectrum.

Eye movements were recorded with an EyeLink1000 Desktop Mount (SR Research) eye tracker, with a sampling rate of 1000 Hz and an eye position resolution of a 20-s arc. All recordings and calibration were binocular. Only right eye data was used for the analyses. Eye movement data from 60 participants reading 120 sentences resulted in a total of 54901 fixations for the controls and 67315 for the people with AD. This data was cleaned from blinks and track losses. For control patients, these procedures resulted in a loss of 1852 records (12%) of data in highly predictable sentences and 7273 (18%) in low predictable sentences. For the AD group, these procedures resulted in a loss of 4676 records (19%) of data in highly predictable sentences and 6618 (15%) in low predictable sentences. Prior to removing the analysis fixations shorter than 51 ms and longer than 750 ms, and fixations on the first and last word of each sentence (see (Kliegl et al., 2006) for a description of the analytic procedure), we measured, for each patient, the elapsed time between the instant when sentences were first presented, and the instant when participants looked at the final spot: mean reading time in high-predictable sentences was 3495 ms vs. 5635 ms (Controls and AD) and 4828 ms vs. 6881 ms (controls and AD) in low predictable-sentences.

2.4 Procedures

Participants gaze was calibrated with a standard 13-point grid for both eyes. After validating the calibration, a trial began with the appearance of a fixation point on the position where the first letter of the sentence was to be presented. As soon as both eyes were detected within a radius equal to 1° from the fixation spot, the sentence was presented. After reading it, participants looked at a dot in the lower

right corner of the screen; when the gaze was detected on the final spot, the trial ended. Occasionally, external factors such as minor movements and slippages of the head gear could cause small drifts. To avoid them, we performed a drift correction before the presentation of each spot.

To assess whether subjects comprehended the texts, they were presented with a three alternative multiple-choice question about the sentence in progress in 20% of the sentence trials. Participants answered the questions by moving a mouse and choosing the response with a mouse click. Overall mean accuracy was 95% ($SD = 3.2\%$) in Control and 91% ($SD = 5.4\%$) in AD. A one-way ANOVA showed no significant differences between comprehension of the answers in Controls and in AD patients. The latter were only marginally less accurate than control subjects, probably because they were in an early stage of the pathology, as indicated by the MMSE and ACE-R values. Once the comprehension test ended, the next trial started with the presentation of the fixation spot. An extra calibration was done after 15 sentences or if the eye tracker did not detect the eye at the initial fixation point within 2 s.

2.5 Sentence corpus

The sentence corpus was composed of short sentences of a line with 75 low predictable sentences, 45 high-predictable sentences and 64 proverbs (e.g., "Maria is always laughing and in a good mood", "Is worth to think before talking" and "A bird in the hand is worth two in the bush")Fernández et al. (2014b). All the sentences comprised a well-balanced number of content and function words and had similar grammatical structure.

2.5.1 Word and Sentence Lengths

Sentences ranged from a minimum of 5 words to a maximum of 14 words. Mean sentence length was 8.1 ($SD = 1.4$) words for low predictability sentences, 7.6 words ($SD = 1.5$) for high predictability sentences and 7.3 words ($SD = 1.9$) for proverbs. Words ranged from 1 to 14 letters. Mean word length was 4.6, 4.1 and 4.0 ($SD = 2.5$, $SD = 2.3$ and $SD = 2.0$) for low-, high-predictable sentences and proverbs, respectively.

2.5.2 Word Frequencies

We used the Spanish Lexical Lexesp corpus Sebastián-Gallés (2000) for assigning a frequency to each word of the sentence corpus. Word frequencies ranged from 1 to 264721 per million, so we transformed it to $log_{10}(frequency)$. Mean $log_{10}(frequency)$ was 3.4 ($SD = 1.3$) for low predictability sentences, 3.4 ($SD = 1.5$) for high predictability sentences and 3.47 ($SD = 1.36$) for proverbs.

2.5.3 Word Predictability

It was measured in an independent experiment with 18 researchers of the Electrical Engineering and Computer Science Department of Universidad Nacional del Sur. We used an incremental cloze task procedure in which participants had to guess the next word given only the prior words of the sentence. Participants guessed the first word of the unknown sentence and entered it via the keyboard. In return, the computer presented the first word of the original sentence on the screen. Responding to this, participants entered their guess for the second word and so on, until a period indicated the end of the sentence. Correct words stayed on the screen. Participants were between 31 and 62 years old and did not participate in the reading experiment. Academic background of the reading experiment group and the cloze task group was similar. Word predictabilities ranged from 0 to 1 with a mean of 0.38 ($SD = 0.36$). The average predictability measured from the cloze task was transformed using a *logit* function $0.5 * ln(pred/(1 - pred))$; predictabilities of zero were replaced with $1/(2 * 18) = -2.55$ and those among the five perfectly predicted words with $(2 * 18 - 1)/(2 * 18) = +2.55$, where 18 represents the number of complete predictability protocols. Mean *logit* predictability was -0.9 ($SD = 0.9$) for low predictability sentences, 0.0 ($SD = 1.29$) for high predictability sentences and 0.08 ($SD = 1.23$) for proverbs.

As in other languages, we find strong correlations in Spanish between word length, word frequency, and word predictability. Long words are of low frequency ($r = -0.80$ and $r = -0.75$ in low and in high predictability sentences, respectively). Frequent words are highly predictable ($r = 0.47$ and $r = 0.37$ in

low and in high predictability sentences, respectively), and highly predictable words tend to be short words ($r = -0.47$ and $r = -0.38$ in low and in high predictability sentences, respectively).

2.6 Used information

The information used for this work was a trial-wise compaction of the original data where we keep descriptors of the mean reading behavior of the subjects in each read sentence. We measured the saccade amplitude, fixation duration and duration of the fixation on a single word of the subject during the reading of each sentence but only kept the mean and the standard deviation. In addition, we measured the total number of fixations and classified them by first pass fixations, refixations, unique fixations and total fixations:

- First pass fixations: The first fixation on a specific word of the sentence.
- Unique fixations: Fixations that occur once in a word that was skipped in the first pass.
- Multiple fixations: Multiple fixations on a word in the first pass.
- Refixations: Fixations that take place once a word that already has a *first pass fixation* or a *unique fixation* implying a regression.

Categorical data as the diagnosis (used as training labels) was replaced by numerical values in order to unify the data types and improve during the classification process. An integer with two possible values was used for the diagnosis information construction. 0 for "Control" and 1 for "AD". The identification and the diagnosis information of the subject were kept apart from the data. A detail of the variables used as input for the model construction is shown in table 1. Since the tag (AD or Control) is associated to the patient and not to each sentence, and, since we use a per-trial classification approach, the subject's tag was applied to all the sentences read by him/her. Following this approach may introduce noise in the classifying stage because, as we use a per-trial classification approach, a control subject could, for example, be distracted during the reading of a specific sentence thus making him read as a non-healthy person. Anyway, the system should be able to detect and ignore these artifacts because many samples are used in the training stage.

Name	**Description**
nw	Number of words in the sentence.
gaze	Global (sentence) mean of the sum of fixation durations on the same word.
sd_gaze	Standard deviation of *gaze*.
as	Mean saccade amplitude in the sentence.
sd_as	Standard deviation of *as*
ntf	Count of the total number of fixations on the sentence.
ntm	Count of the number of multifixations on the sentence.
dfp	Mean duration of the first pass fixations on the sentence.
sd_dfp	Standard deviation of *dfp*.
fpp	Count of the number of first pass fixations on the sentence.
rf	Count of refixations on the sentence.
nfu	Count of unique fixations on the sentence.
dfu	Mean duration of unique fixations on the sentence.
sd_dfu	Standard deviation of *dfu*.

Table 1: Used variables for the model construction.

All the data was previously outlier-checked by establishing a dropout policy in order to use a cleaner dataset. The outlier check policy consisted of finding the mean and the SD of each condition group and checking the two groups separately. All the trials where the standard deviation was bigger than two times the standard deviation of the group were considered as outliers and was dropped out, resulting in a 10% samples lost. The sentence identication, order, and type were kept separate from the training

information because the standard deviations in the information of the AD patients for proverbs and high predictability sentences appear to be particularly high after the data is outlier-checked, causing a highly unbalanced dataset. Finally, 90% of the data was used in the training stage and the remaining 10% for preliminary testing in order to keep a clean training set and avoiding over-tting by ensuring even results in the data used for training and in the data unknown by the system used for validation.

2.7 Deep learning with denoising sparse-autoencoders

In this work we used *sparse-autoencoders* for the codification stage. The sparse-autoencoders work just as regular autoencoders, i.e. there are neural networks under supervised learning with the targets set equal to the input (the identity) but, in the case of sparse-autoencoders, an average number of activations per neuron restriction was applied in the hidden layer by penalizing the average number of activations different from the desired (known as sparsity proportion) adding a penalty term to the cost function. This restriction is introduced so that each neuron specializes on a particular feature. The lower the sparsity proportion, the more specific the feature. The resulting trained neural network can be thought of as: an encoder, involving the input and the hidden layer, and a decoder, involving the hidden and the output layer. In this case, we set an activation restriction equal to 10%.

In a denoising autoencoder, the idea is to force the hidden layer to discover more robust features and to prevent it from simply learning the identity, by training the autoencoder to reconstruct the input from a corrupt version. The altered version of the input was generated by introducing noise, which was obtained by clamping some of the fields to zero. The corrupt data was used as the sparse-autoencoder input, and the clean (unaltered) data as the target. Using this type of data corruption mechanism forces the network to learn a way of reconstructing a field based on others. This, combined with the sparsity restriction, results in more robust features.

The deep learning neural network was built using two stages of these denoising sparse-autoencoders. In each stage, we train the autoencoder by corrupting the encoded clean data obtained from the previous stage, and providing it to the next as its input. At the end of the two stages, we set a softmax layer as a classifier by training it with uncorrupted data and the corresponding tag. As we used a per-trial classification approach, the patient diagnosis was extended to all the sentences read by him and the classifier was trained with this data as the target. The softmax layer is a non-linear, multiclass generalization of the binary Logistic Regression, and its output is the "probability" for each class (we quoted the word "probability" because it's shape depends on of the regularization used in the training stage, it can be more diffuse or peaky).

3 Results

Several configurations were generated by varying the sparsity proportion, the number of units and layers and the shape of the network (same vs. decreasing number of units between layers). We adopted the one that produced the best results which consisted of two layers of denoising sparse-autoencoders with 16 and 4 hidden units in the hidden layer each, using a sparsity proportion of 10%. After the training of the network, a series of tests were performed with data not included in the training dataset. This dataset was composed of 462 sentences randomly chosen from the original dataset was 222 are sentences read by control patients and 240 sentences read by AD patients. As mentioned above, we used a softmax layer for the classification. It was trained using the condition translation with 0 for control subjects and 1 for AD subjects. This means that we have a single class "AD" and, since the output of the classier is a real number between 0 and 1, the read sentences classified by the network with values close to 0 have a small "probability" of being read by an AD patient (i.e. high probability of being read by a control patient) and vice-versa. The "ground truth" values are known, so we can split the output into groups and observe the number of sentences misclassified by the network. Based on this, we show the output of the network where values below 0.5 are considered as classified as Control, and higher values are considered as classified as AD (see Figure 1). As can be seen, the output of the network was consistent with the expected values.

Now, we can round the values so we can plot a confusion matrix and approximate the number of misclassified sentences in order to measure the performance of the network. In the figure 2 we can see a confusion matrix of the output. The "X" axis represents the expected output values and the "Y"

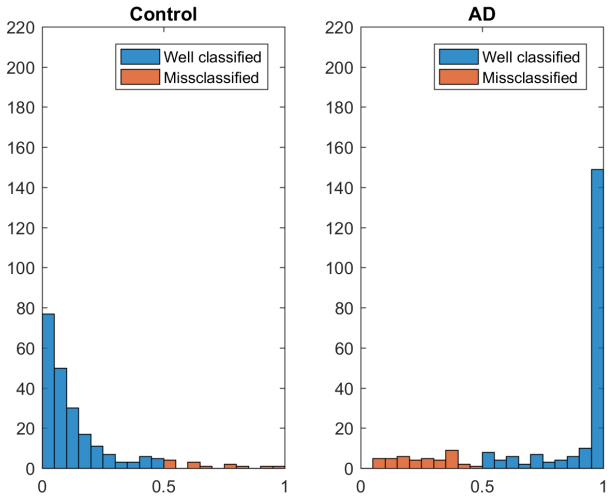

Figure 1: Classification results histogram split by "ground truth" values. Values below 0.5 are considered as classified as Control, and higher values are considered as classified as AD.

axis the rounded output of the network. As can be seen, the overall performance of the network was good giving a 88.3% of well-classified sentences. The performance of the network using sentences read by control patients (94.1% correctness) approximates to the performance observed using sentences read by AD patients (82.6% correctness).

In the other hand, misclassified sentences were not concentrated on a particular type of sentence as can be seen in figure 3. Here we can see the original concentration of sentences types in the testing dataset and the correctness of the classification following the mentioned method.

This result, added to the fact that neither was a concentration of misclassication in a particular sentence, may suggest that most of the misclassications occurred due to presumably stochastic processes. In addition, the trained networks were evaluated using a spread result test in order to determine the softness of the model. These tests checked if similar information is encoded in a similar way in the subsequent stages of the network. A significant differentiation in later stages of encoding may show overfitting in the network (and/or in the different stages).

Two subsets (one composed of AD patients and the other of control subjects) of trials that have similar values for the input in each field are shown in Figure 4. As shown, similar input values map to similar encoded values on each stage of the autoencoders. This is because the modeled function is smooth. Furthermore, the data through the subsequent stages of codification tends to group. These results have shown that the output information such as the encoding in the different stages are reliable. On the other hand, they show that certain neurons in later stages tend to specialize on the detection of specific AD or Control input features.

4 Conclussions and future work

The results showed that using a deep-learning architecture for identifying the characteristic eye movements patterns of neurodegenerative diseases like Alzheimer's disease is a good approach since this technology is focused on pattern finding and is suitable for this work. Moreover, the high performance in a per-trial classification approach, leads us to conclude that, since a single patient reads many sentences, the assertion rate is bigger than the 88.3% reported in this work. Additionally, we asked the head psychiatrists of some AD patients (that were not included in training process) to elaborate a score of the overall performance of each patient on the traditional tests using a scale from 0 to 1, without knowing the results given by our network. The process of creating the score required that the physicians have a deep knowledge of the psychiatric history of each patient, the recompilation of the results of every neuropsychological tests made by the patient and its comprehension. Table 2 shows the scores given by

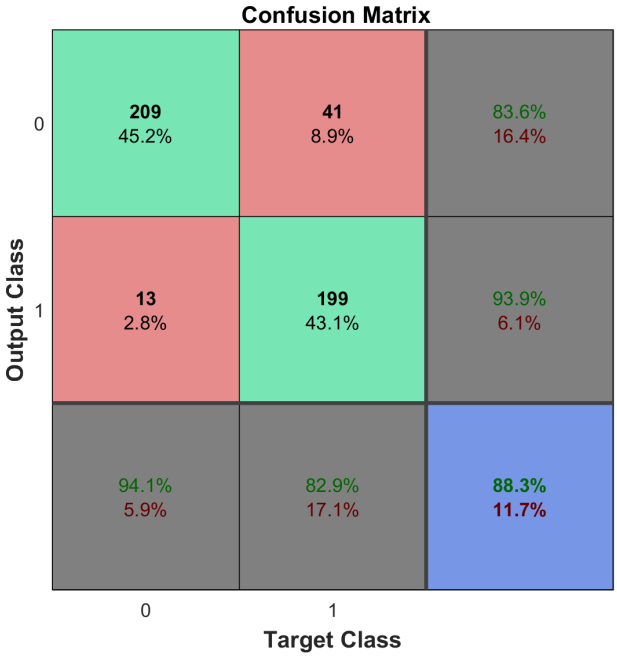

Figure 2: Classification results rounded.

the psychiatrists compared with the mean value obtained in our network for all the sentences read by the subjects and its SD. As seen, for most of the patients, the values obtained are very similar to the scores given by the psychiatrists with a mean value of 0.19 ($SD = 0.15$). The obtained results show that the created marker is reasonably close to the score but involves a much simpler process.

Finding a better way to interpret the output of the classifier is left as future work. An improvement is required because values near 0.5 are not determined to be identified as AD or Control (equal "probabilities"). The policy chosen in this work is a first rough approach that doesnt reflect the actual power of the network. Using a fuzzy-logic encoder to obtain the overall diagnosis of a patient might lead to a more

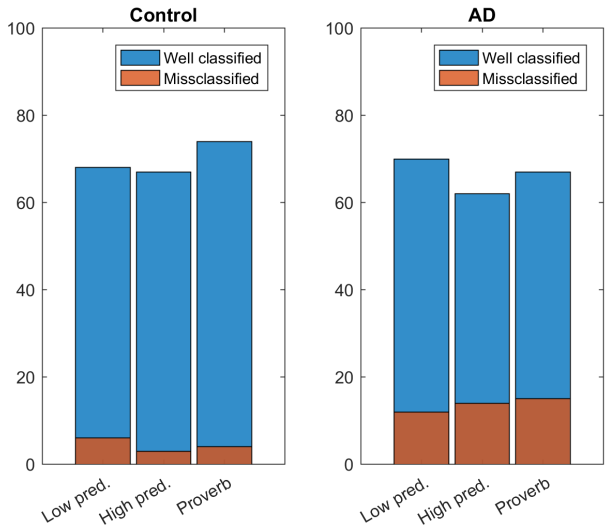

Figure 3: Concentration of misclassified sentences by type.

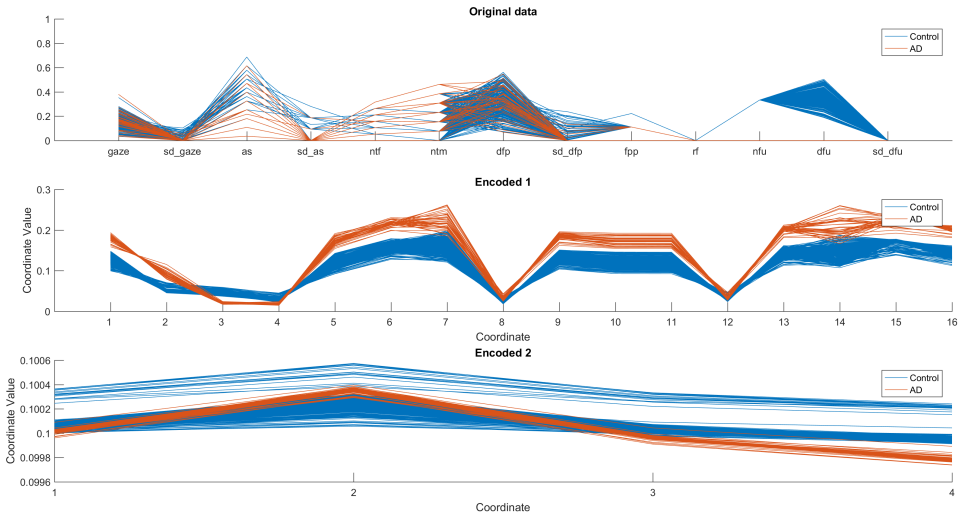

Figure 4: Two subsets (one composed of AD patients and the other of control subjects) of trials that have similar values for the input in each field and its codification during the different stages of the network showing similar values in each codification stage.

accurate result. Determining whether the number given by the classier is related to the severity of the disease is left as a future improvement. This task is particularly difficult since there are no ground truth measurements to corroborate the information due to the current psychological testing methods. Anyway, although we adopted a per-trial classification approach, it's easy to think that the overall diagnosis may be related to a measurement extracted from the entire test and not from a single trial as shown before. Besides this, here, as a validation test, we used the mean of all the reported results.

IDPat	**Mean**	**SD**	**Score**	**Difference**
58	0,97	0,17	0,9	0,07
57	0,95	0,17	0,5	0,45
66	0,49	0,32	0,5	0,01
60	0,95	0,16	0,6	0,35
56	0,96	0,16	0,8	0,16
55	0,94	0,17	0,7	0,24
63	0,87	0,25	0,8	0,07
64	0,51	0,34	0,5	0,01
70	0,90	0,24	0,5	0,40
69	0,84	0,25	0,5	0,34
65	0,91	0,22	0,6	0,31
59	0,58	0,36	0,6	0,02
71	0,75	0,33	0,6	0,15
62	0,76	0,32	0,6	0,16
67	0,47	0,35	0,5	0,03
68	0,40	0,31	0,8	0,40
53	0,84	0,31	0,8	0,04
			Mean	0,19
			SD	0,15

Table 2: Comparison of mean value given by the network and performance score given by psychiatrists.

References

Li Deng and Dong Yu. Deep learning: Methods and applications. *Foundations and Trends in Signal Processing*, 7(3–4):197–387, 2014.

Gerardo Fernández, Pablo Mandolesi, Nora P Rotstein, Oscar Colombo, Osvaldo Agamennoni, and Luis E Politi. Eye movement alterations during reading in patients with early alzheimer disease, eye movement behavior in alzheimer disease patients. *Investigative ophthalmology & visual science*, 54 (13):8345–8352, 2013.

Gerardo Fernández, Jochen Laubrock, Pablo Mandolesi, Oscar Colombo, and Osvaldo Agamennoni. Registering eye movements during reading in alzheimers disease: Difficulties in predicting upcoming words. *Journal of clinical and experimental neuropsychology*, 36(3):302–316, 2014a.

Gerardo Fernández, Diego E Shalom, Reinhold Kliegl, and Mariano Sigman. Eye movements during reading proverbs and regular sentences: The incoming word predictability effect. *Language, Cognition and Neuroscience*, 29(3):260–273, 2014b.

Gerardo Fernández, Liliana R Castro, Marcela Schumacher, and Osvaldo E Agamennoni. Diagnosis of mild alzheimer disease through the analysis of eye movements during reading. *Journal of integrative neuroscience*, 14(01):121–133, 2015a.

Gerardo Fernández, Marcela Schumacher, Liliana Castro, David Orozco, and Osvaldo Agamennoni. Patients with mild alzheimers disease produced shorter outgoing saccades when reading sentences. *Psychiatry research*, 229(1):470–478, 2015b.

Gerardo Fernández, Salvador Guinjoan, Marcelo Sapognikoff, David Orozco, and Osvaldo Agamennoni. Contextual predictability enhances reading performance in patients with schizophrenia. *Psychiatry research*, 241:333–339, 2016a.

Gerardo Fernández, Facundo Manes, Luis E Politi, David Orozco, Marcela Schumacher, Liliana Castro, Osvaldo Agamennoni, and Nora P Rotstein. Patients with mild alzheimers disease fail when using their working memory: Evidence from the eye tracking technique. *Journal of Alzheimer's Disease*, 50 (3):827–838, 2016b.

Gerardo Fernández, Marcelo Sapognikoff, Salvador Guinjoan, David Orozco, and Osvaldo Agamennoni. Word processing during reading sentences in patients with schizophrenia: evidences from the eyetracking technique. *Comprehensive psychiatry*, 68:193–200, 2016c.

Marshal F Folstein, Susan E Folstein, and Paul R McHugh. mini-mental state: a practical method for grading the cognitive state of patients for the clinician. *Journal of psychiatric research*, 12(3):189–198, 1975.

Philip S Holzman, Leonard R Proctor, Deborah L Levy, Nicholas J Yasillo, Herbert Y Meltzer, and Stephen W Hurt. Eye-tracking dysfunctions in schizophrenic patients and their relatives. *Archives of general psychiatry*, 31(2):143–151, 1974.

William G Iacono, Margaret Moreau, Morton Beiser, Jonathan AE Fleming, and Tsung-Yi Lin. Smooth-pursuit eye tracking in first-episode psychotic patients and their relatives. *Journal of Abnormal Psychology*, 101(1):104, 1992.

Jennifer L Kellough, Christopher G Beevers, Alissa J Ellis, and Tony T Wells. Time course of selective attention in clinically depressed young adults: An eye tracking study. *Behaviour research and therapy*, 46(11):1238–1243, 2008.

Reinhold Kliegl, Antje Nuthmann, and Ralf Engbert. Tracking the mind during reading: the influence of past, present, and future words on fixation durations. *Journal of experimental psychology: General*, 135(1):12, 2006.

Rongjian Li, Wenlu Zhang, Heung-Il Suk, Li Wang, Jiang Li, Dinggang Shen, and Shuiwang Ji. Deep learning based imaging data completion for improved brain disease diagnosis. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*, pages 305–312. Springer, 2014.

Siqi Liu, Sidong Liu, Weidong Cai, Sonia Pujol, Ron Kikinis, and Dagan Feng. Early diagnosis of alzheimer's disease with deep learning. In *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on*, pages 1015–1018. IEEE, 2014.

Guy McKhann, David Drachman, Marshall Folstein, Robert Katzman, Donald Price, and Emanuel M Stadlan. Clinical diagnosis of alzheimer's disease report of the nincds-adrda work group* under the auspices of department of health and human services task force on alzheimer's disease. *Neurology*, 34 (7):939–939, 1984.

Eneida Mioshi, Kate Dawson, Joanna Mitchell, Robert Arnold, and John R Hodges. The addenbrooke's cognitive examination revised (ace-r): a brief cognitive test battery for dementia screening. *International journal of geriatric psychiatry*, 21(11):1078–1085, 2006.

Keith Rayner. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124(3):372, 1998.

Deborah Riby and Peter JB Hancock. Looking at movies and cartoons: eye-tracking evidence from williams syndrome and autism. *Journal of Intellectual Disability Research*, 53(2):169–181, 2009.

Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

Núria Sebastián-Gallés. *LEXESP: Léxico informatizado del español*. Edicions Universitat Barcelona, 2000.

Heung-Il Suk and Dinggang Shen. Deep learning-based feature representation for ad/mci classification. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pages 583–590. Springer, 2013.

Heung-Il Suk, Seong-Whan Lee, Dinggang Shen, Alzheimer's Disease Neuroimaging Initiative, et al. Hierarchical feature representation and multimodal fusion with deep learning for ad/mci diagnosis. *NeuroImage*, 101:569–582, 2014.

G Waldemar, B Dubois, M Emre, J Georges, IG McKeith, M Rossor, P Scheltens, P Tariska, and B Winblad. Recommendations for the diagnosis and management of alzheimer's disease and other disorders associated with dementia: Efns guideline. *European Journal of Neurology*, 14(1):e1–e26, 2007.